arXiv:1608.05995v5 [stat.ML] 25 Oct 2016

A Non-convex One-Pass Framework for Generalized Factorization Machine and Rank-One Matrix Sensing

Ming Lin*[1] and Jieping Ye†[1]

[1]Department of Computational Medicine and Bioinformatic,
University of Michigan, Ann Arbor, MI 48109

August 21, 2016

Abstract

We develop an efficient alternating framework for learning a generalized version of Factorization Machine (gFM) on steaming data with provable guarantees. When the instances are sampled from d dimensional random Gaussian vectors and the target second order coefficient matrix in gFM is of rank k, our algorithm converges linearly, achieves $O(\epsilon)$ recovery error after retrieving $O(k^3 d \log(1/\epsilon))$ training instances, consumes $O(kd)$ memory in one-pass of dataset and only requires matrix-vector product operations in each iteration. The key ingredient of our framework is a construction of an estimation sequence endowed with a so-called Conditionally Independent RIP condition (CI-RIP). As special cases of gFM, our framework can be applied to symmetric or asymmetric rank-one matrix sensing problems, such as inductive matrix completion and phase retrieval.

1 Introduction

Linear models are one of the foundations of modern machine learning due to their strong learning guarantees and efficient solvers [Koltchinskii, 2011]. Conventionally linear models only consider the first order information of the input feature which limits their capacity in non-linear problems. Among various efforts extending linear models to the non-linear domain, the Factorization Machine [Rendle, 2010] (FM) captures the second order information by modeling the pairwise feature interaction in regression under low-rank constraints. FMs have

*linmin@umich.edu
†jpye@umich.edu

been found successful in many applications, such as recommendation systems [Rendle et al., 2011] and text retrieval [Hong et al., 2013]. In this paper, we consider a generalized version of FM called gFM which removes several redundant constraints in the original FM such as positive semi-definite and zero-diagonal, leading to a more general model without sacrificing its learning ability. From theoretical side, the gFM includes rank-one matrix sensing [Zhong et al., 2015, Chen et al., 2015, Cai and Zhang, 2015, Kueng et al., 2014] as a special case, where the latter one has been studied widely in context such as inductive matrix completion [Jain and Dhillon, 2013] and phase retrieval [Candes et al., 2011].

Despite of the popularity of FMs in industry, there is rare theoretical study of learning guarantees for FMs. One of the main challenges in developing a provable FM algorithm is to handle its symmetric rank-one matrix sensing operator. For conventional matrix sensing problems where the matrix sensing operator is RIP, there are several alternating methods with provable guarantees [Hardt, 2013, Jain et al., 2013, Hardt and Wootters, 2014, Zhao et al., 2015a,b]. However, for a symmetric rank-one matrix sensing operator, the RIP condition doesn't hold trivially which turns out to be the main difficulty in designing efficient provable FM solvers.

In rank-one matrix sensing, when the sensing operator is asymmetric, the problem is also known as inductive matrix completion which can be solved via alternating minimization with a global linear convergence rate [Jain and Dhillon, 2013, Zhong et al., 2015]. For symmetric rank-one matrix sensing operators, we are not aware of any efficient solver by the time of writing this paper. In a special case when the target matrix is of rank one, the problem is called "phase retrieval" whose convex solver is first proposed by Candes et al. [2011] then alternating methods are provided in [Lee et al., 2013, Netrapalli et al., 2013]. While the target matrix is of rank $k > 1$, only convex methods minimizing the trace norm have been proposed recently, which are computationally expensive [Kueng et al., 2014, Cai and Zhang, 2015, Chen et al., 2015, Davenport and Romberg, 2016].

Despite of the above fundamental challenges, extending rank-one matrix sensing algorithm to gFM itself is difficult. Please refer to Section 2.1 for an in-depth discussion. The main difficulty is due to the first order term in the gFM formulation, which cannot be trivially converted to a standard matrix sensing problem.

In this paper, we develop a unified theoretical framework and an efficient solver for generalized Factorization Machine and its special cases such as rank-one matrix sensing, either symmetric or asymmetric. The key ingredient is to show that the sensing operator in gFM satisfies a so-called Conditionally Independent RIP condition (CI-RIP, see Definition 2) . Then we can construct an estimation sequence via noisy power iteration [Hardt and Price, 2013]. Unlike previous approaches, our method does not require alternating minimization or choosing the step-size as in alternating gradient descent. The proposed method works on steaming data, converges linearly and has $O(kd)$ space complexity for a d-dimension rank-k gFM model. The solver achieves $O(\epsilon)$ recovery error after retrieving $O(k^3 d \log(1/\epsilon))$ training instances.

The remainder of this paper is organized as following. In Section 2, we introduce necessary notation and background of gFM. Subsection 2.1 investigates several fundamental challenges in depth. Section 3 presents our algorithm, called One-Pass gFM, followed by its theoretical guarantees. Our analysis framework is presented in Section 4. Section 5 concludes this paper.

2 Generalized Factorization Machine (gFM)

In this section, we first introduce necessary notation and background of FM and its generalized version gFM. Then in Subsection 2.1, we reveal the connection between gFM and rank-one matrix sensing followed by several fundamental challenges encountered when applying frameworks of rank-one matrix sensing to gFM.

The FM predicts the labels of instances by not only their features but also high order interactions between features. In the following, we focus on the second order FM due to its popularity. Suppose we are given N training instances $\boldsymbol{x}_i \in \mathbb{R}^d$ independently and identically (I.I.D.) sampled from the standard Gaussian distribution and so are their associated labels $y_i \in \mathbb{R}$. Denote the feature matrix $X = [\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_n] \in \mathbb{R}^{d\times n}$ and the label vector $\boldsymbol{y} = [y_1, y_2, \cdots, y_n]^\top \in \mathbb{R}^n$. In second order FM, y_i is assumed to be generated from a target vector $\boldsymbol{w}^* \in \mathbb{R}^d$ and a target rank k matrix $M^* \in \mathbb{R}^{d\times d}$ satisfying

$$y_i = \boldsymbol{x}_i{}^\top \boldsymbol{w}^* + \boldsymbol{x}_i{}^\top M^* \boldsymbol{x}_i + \xi_i \tag{1}$$

where ξ_i is a random subgaussian noise with proxy variance ξ^2 . It is often more convenient to write Eq. (1) in matrix form. Denote the linear operator $\mathcal{A} : \mathbb{R}^{d\times d} \to \mathbb{R}^n$ as $\mathcal{A}(M) \triangleq [\langle A_1, M\rangle, \langle A_2, M\rangle, \cdots, \langle A_n, M\rangle]^\top$ where $A_i = \boldsymbol{x}_i \boldsymbol{x}_i{}^\top$. Then Eq. (1) has a compact form:

$$\boldsymbol{y} = X^\top \boldsymbol{w}^* + \mathcal{A}(M^*) + \boldsymbol{\xi} \ . \tag{2}$$

The FM model given by Eq. (2) consists of two components: the first order component $X^\top \boldsymbol{w}^*$ and the second order component $\mathcal{A}(M^*)$. The component $\mathcal{A}(M^*)$ is a symmetric rank-one Gaussian measurement since $\mathcal{A}_i(M) = \boldsymbol{x}_i{}^\top M \boldsymbol{x}_i$ where the left/right design vectors ($\boldsymbol{x}_i$ and $\boldsymbol{x}_i{}^\top$) are identical. The original FM requires that M^* should be positive semi-definite and the diagonal elements of M^* should be zero. However our analysis shows that both constraints are redundant for learning Eq. 2. Therefore in this paper we consider a generalized version of FM which we call gFM where M^* is only required to be symmetric and low rank. To make the recovery of M^* well defined, it is necessary to assume M^* to be symmetric. Indeed for any asymmetric matrix M^*, there is always a symmetric matrix $M^*_{\text{sym}} = (M^* + M^{*\top})/2$ such that $\mathcal{A}(M^*) = \mathcal{A}(M^*_{\text{sym}})$ thus the symmetric constraint does not affect the model. Another standard assumption in rank-one matrix sensing is that the rank of M^* should be no more than k for $k \ll d$. When $\boldsymbol{w}^* = 0$, gFM is equal to the symmetric rank-one matrix sensing problem. Recent researches have proposed several convex programming methods based on the trace norm minimization to recover M^* with

a sampling complexity on order of $O(k^3 d)$ [Candes et al., 2011, Cai and Zhang, 2015, Kueng et al., 2014, Chen et al., 2015, Zhong et al., 2015]. Some authors also call gFM as second order polynomial network [Blondel et al., 2016].

When d is much larger than k, the convex programming on the trace norm or nuclear norm of M^* becomes difficult since M^* can be a $d \times d$ dense matrix. Although modern convex solvers can scale to large d with reasonable computational cost, a more popular strategy to efficiently estimate $\boldsymbol{w}^*$ and M^* is to decompose M^* as $UV^\top$ for some $U, V \in \mathbb{R}^{d\times k}$, then alternatively update $\boldsymbol{w}$, U, V to minimize the empirical loss function

$$\min_{\boldsymbol{w},U,V} \frac{1}{2N}\|\boldsymbol{y} - X^\top \boldsymbol{w} - \mathcal{A}(UV^\top)\|_2^2 \,. \tag{3}$$

The loss function in Eq. (3) is non-convex. It is even unclear whether an estimator of the optimal solution $\{\boldsymbol{w}^*, M^*\}$ of Eq. (3) with a polynomial time complexity exists or not.

In our analysis, we denote $M + O(\epsilon)$ as a matrix M plus a perturbation matrix whose spectral norm is bounded by ϵ. We use $\|\cdot\|_2$, $\|\cdot\|_F$, $\|\cdot\|_*$ to denote the matrix spectral norm, Frobenius norm and nuclear norm respectively. To abbreviate the high probability bound, we denote $C = \mathrm{polylog(d, n, T, 1/\eta)}$ to be a constant polynomial logarithmic in $\{d, n, T, 1/\eta\}$. The eigenvalue decomposition of M^* is $M^* = U^*\Lambda^* U^{*\top}$ where $U^* \in \mathbb{R}^{d\times k}$ is the top-k eigenvectors of M^* and $\Lambda^* = \mathrm{diag}(\lambda_1^*, \lambda_2^*, \cdots, \lambda_k^*)$ are the corresponding eigenvalues sorted by $|\lambda_i| \geq |\lambda_{i+1}|$. Let $\sigma_i^* = |\lambda_i^*|$ denote the singular value of M^* and $\sigma_i\{M\}$ be the i-th largest singular value of M. $U_\perp^*$ denotes an matrix whose columns are the orthogonal basis of the complementary subspace of U^*.

2.1 gFM and Rank-One Matrix Sensing

When $\boldsymbol{w}^* = 0$ in Eq. (1), the gFM becomes the symmetric rank-one matrix sensing problem. While the recovery ability of rank-one matrix sensing is somehow provable recently despite of the computational issue, it is not the case for gFM. It is therefore important to discuss the differences between gFM and rank-one matrix sensing to give us a better understanding of the fundamental barriers in developing provable gFM algorithm.

In the rank-one matrix sensing problem, a relaxed setting is to assume that the sensing operator is asymmetric, which is defined by $\mathcal{A}_i^{\mathrm{asy}}(M) = \boldsymbol{u}_i{}^\top M \boldsymbol{v}_i$ where $\boldsymbol{u}_i$ and $\boldsymbol{v_i}$ are independent random vectors. Under this setting, the recovery ability of alternating methods is provable [Jain and Dhillon, 2013]. However, existing analyses cannot be generalized to their symmetric counterpart, since $\boldsymbol{u}_i$ and $\boldsymbol{v}_i$ are not allowed to be dependent in these frameworks. For example, the sensing operator $\mathcal{A}^{\mathrm{asy}}(\cdot)$ is unbiased ($E\mathcal{A}^{\mathrm{asy}}(\cdot) = 0$) but the symmetric sensing operator is clearly not [Cai and Zhang, 2015]. Therefore, the asymmetric setting oversimplifies the problem and loses important structure information which is critical to gFM.

As for the symmetric rank-one matrix sensing operator, the state-of-the-art estimator is based on the trace norm convex optimization [Tropp, 2014,

Chen et al., 2015, Cai and Zhang, 2015], which is computationally expensive. When $\boldsymbol{w}^* \neq \mathbf{0}$, the gFM has an extra perturbation term $X^\top \boldsymbol{w}^*$. This first order perturbation term turns out to be a fundamental challenge in theoretical analysis. One might attempt to merge $\boldsymbol{w}^*$ into M^* in order to convert gFM as a rank $(k+1)$ matrix sensing problem. For example, one may extend the feature $\hat{\boldsymbol{x}}_i \triangleq [\boldsymbol{x}_i, 1]^\top$ and the matrix $\hat{M}^* = [M^*; \boldsymbol{w}^{*\top}] \in \mathbb{R}^{(d+1)\times d}$. However, after this simple extension, the sensing operator becomes $\hat{\mathcal{A}}(M^*) = \hat{\boldsymbol{x}}_i{}^\top \hat{M}^* \boldsymbol{x}_i$. It is no longer symmetric. The left/right design vector is neither independent nor identical. Especially, not all dimensions of $\hat{\boldsymbol{x}}_i$ are random variables. According to the above discussion, the conditions to guarantee the success of rank-one matrix sensing do not hold after feature extension and all the mentioned analyses cannot be directly applied.

3 One-Pass gFM

In this section, we present the proposed algorithm, called One-Pass gFM followed by its theoretical guarantees. We will focus on the intuition of our algorithm. A rigorous theoretical analysis is presented in the next section.

The One-Pass gFM is a mini-batch algorithm. In each mini-batch, it processes n training instances and then alternatively updates parameters. The iteration will continue until T mini-batch updates. Since gFM deals with a non-convex learning problem, the conventional gradient descent framework hardly works to show the global convergence. Instead, our method is based on a construction of an estimation sequence. Intuitively, when $\boldsymbol{w}^* = \mathbf{0}$, we will show in the next section that $\frac{1}{n}\mathcal{A}'\mathcal{A}(M) \approx 2M + \mathrm{tr}(M)I$ and $\mathrm{tr}(M) \approx \frac{1}{n}\mathbf{1}^\top \mathcal{A}(M)$. Since $\boldsymbol{y} \approx \mathcal{A}(M^*)$, we can estimate M^* via $\frac{1}{2n}\mathcal{A}'(\boldsymbol{y}) - \frac{1}{n}\mathbf{1}^\top \boldsymbol{y} I$. But this simple construction cannot generate a convergent estimation sequence since the perturbation terms in the above approximate equalities cannot be reduced along iterations. To overcome this problem, we replace $\mathcal{A}(M^*)$ with $\mathcal{A}(M^* - M^{(t)})$ in our construction. Then the perturbation terms will be on order of $O(\|M^* - M^{(t)}\|_2)$. When $\boldsymbol{w}^* \neq \mathbf{0}$, we can apply a similar trick to construct its estimation sequence via the second and the third order moments of X. Algorithm 1 gives a step-by-step description of our algorithm[1].

In Algorithm 1, we only need to store $\boldsymbol{w}^{(t)} \in \mathbb{R}^d$, $U^{(t)}, V^{(t)} \in \mathbb{R}^{d\times k}$. Therefore the space complexity is $O(d+kd)$. The auxiliary variables $M^{(t)}, H_1^{(t)}, h_2^{(t)}, \boldsymbol{h}_3^{(t)}$ can be implicitly presented by $\boldsymbol{w}^{(t)}, U^{(t)}, V^{(t)}$. In each mini-batch updating, we only need matrix-vector product operations which can be efficiently implemented on many computation architectures. We use truncated SVD to initialize gFM, a standard initialization step in matrix sensing. We do not require this step to be computed exactly but up to an accuracy of $O(\delta)$ where δ is the RIP constant. The QR step on line 6 requires $O(k^2 d)$ operations. Compared with SVD which requires $O(kd^2)$ operations, the QR step is much more efficient when $d \gg k$. Algorithm 1 retrieves instances streamingly, a favorable behavior

[1]Implementation is available from `https://minglin-home.github.io/`

Algorithm 1 One-Pass gFM

Require: The mini-batch size n, number of total mini-batch update T, training instances $X = [\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots \boldsymbol{x}_{nT}]$, $\boldsymbol{y} = [y_1, y_2, \cdots, y_{nT}]^\top$, desired rank $k \geq 1$.

Ensure: $\boldsymbol{w}^{(T)}, U^{(T)}, V^{(T)}$.

1: Define $M^{(t)} \triangleq (U^{(t)}V^{(t)\top} + V^{(t)}U^{(t)\top})/2$, $H_1^{(t)} \triangleq \frac{1}{2n}\mathcal{A}'(\boldsymbol{y} - \mathcal{A}(M^{(t)}) - X^{(t)\top}\boldsymbol{w}^{(t)})$, $h_2^{(t)} \triangleq \frac{1}{n}\mathbf{1}^\top(\boldsymbol{y} - \mathcal{A}(M^{(t)}) - X^{(t)\top}\boldsymbol{w}^{(t)})$, $\boldsymbol{h}_3^{(t)} \triangleq \frac{1}{n}X^{(t)}(\boldsymbol{y} - \mathcal{A}(M^{(t)}) - X^{(t)\top}\boldsymbol{w}^{(t)})$.

2: Initialize: $\boldsymbol{w}^{(0)} = \mathbf{0}$, $V^{(0)} = 0$. $U^{(0)} = \mathrm{SVD}(H_1^{(0)} - \frac{1}{2}h_2^{(0)}I, k)$, that is, the top-$k$ left singular vectors.

3: **for** $t = 1, 2, \cdots, T$ **do**

4: Retrieve n training instances $X^{(t)} = [\boldsymbol{x}_{(t-1)n+1}, \cdots, \boldsymbol{x}_{(t-1)n+n}]$. Define $\mathcal{A}(M) \triangleq [X_i^{(t)\top} M X_i^{(t)}]_{i=1}^n$.

5: $\hat{U}^{(t)} = (H_1^{(t-1)} - \frac{1}{2}h_2^{(t-1)}I + M^{(t-1)\top})U^{(t-1)}$.

6: Orthogonalize $\hat{U}^{(t)}$ via QR decomposition: $U^{(t)} = \mathrm{QR}\left(\hat{U}^{(t)}\right)$.

7: $\boldsymbol{w}^{(t)} = \boldsymbol{h}_3^{(t-1)} + \boldsymbol{w}^{(t-1)}$.

8: $V^{(t)} = (H_1^{(t-1)} - \frac{1}{2}h_2^{(t-1)}I + M^{(t-1)})U^{(t)}$

9: **end for**

10: **Output:** $\boldsymbol{w}^{(T)}, U^{(T)}, V^{(T)}$.

on systems with high speed cache. Finally, we export $\boldsymbol{w}^{(T)}, U^{(T)}, V^{(T)}$ as our estimation of $\boldsymbol{w}^* \approx \boldsymbol{w}^{(T)}$ and $M^* \approx U^{(T)}V^{(T)\top}$.

Our main theoretical result is presented in the following theorem, which gives the convergence rate of recovery and sampling complexity of gFM when M^* is low rank and the noise $\boldsymbol{\xi} = \mathbf{0}$.

Theorem 1. *Suppose $\boldsymbol{x}_i$'s are independently sampled from the standard Gaussian distribution. M^* is a rank k matrix. The noise $\boldsymbol{\xi} = \mathbf{0}$. Then with a probability at least $1 - \eta$, there exists a constant C and a constant $\delta < 1$ such that*

$$\|\boldsymbol{w}^* - \boldsymbol{w}^{(t)}\|_2 + \|M^* - M^{(t)}\|_2 \leq \delta^t(\|\boldsymbol{w}^*\|_2 + \|M^*\|_2)$$

provided $n \geq C(4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)^2 k^3 d/\delta^2$, $\delta \leq \frac{(4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)\sigma_k^*}{4\sqrt{5}\sigma_1^* + 3\sigma_k^* + 4\sqrt{5}\|\boldsymbol{w}^*\|_2^2}$.

Theorem 1 shows that $\{\boldsymbol{w}^{(t)}, M^{(t)}\}$ will converge to $\{\boldsymbol{w}^*, M^*\}$ linearly. The convergence rate is controlled by δ, whose value is on order of $O(1/\sqrt{n})$. A small δ will result in a fast convergence rate but a large sampling complexity. To reduce the sampling complexity, a large δ is preferred. The largest allowed δ is bounded by $O(1/(\|M^*\|_2 + \|\boldsymbol{w}^*\|_2))$. The sampling complexity is $O((\sigma_1^*/\sigma_k^*)^2 k^3 d)$. If M^* is not well conditioned, it is possible to remove $(\sigma_1^*/\sigma_k^*)^2$ in the sampling complexity by a procedure called "soft-deflation" [Jain et al., 2013, Hardt and Wootters, 2014]. By theorem 1, gFM achieves ϵ recovery error after retrieving $nT = O(k^3 d \log((\|\boldsymbol{w}^*\|_2 + \|M^*\|_2)/\epsilon))$ instances.

The noisy case where M^* is not exactly low rank and $\xi > 0$ is more intricate therefore we postpone it to Subsection 4.1. The main conclusion is similar to the noise-free case Theorem 1 under a small noise assumption.

4 Theoretical Analysis

In this section, we give the sketch of our proof of Theorem 1. Omitted details are postponed to appendix.

From high level, our proof constructs an estimation sequence $\{\widetilde{\boldsymbol{w}}^{(t)}, \widetilde{M}^{(t)}, \epsilon_t\}$ such that $\epsilon_t \to 0$ and $\|\boldsymbol{w}^* - \widetilde{\boldsymbol{w}}^{(t)}\|_2 + \|M^* - \widetilde{M}^{(t)}\|_2 \leq \epsilon_t$. In conventional matrix sensing, this construction is possible when the sensing matrix satisfies the Restricted Isometric Property (RIP) [Candès and Recht, 2009]:

Definition 2 (ℓ_2-norm RIP)**.** A sensing operator $\mathcal{A}$ is ℓ_2-norm δ_k-RIP if for any rank k matrix M,

$$(1-\delta_k)\|M\|_F \leq \frac{1}{n}\|\mathcal{A}(M)\|_2^2 \leq (1+\delta_k)\|M\|_F \ .$$

When $\mathcal{A}$ is ℓ_2-norm δ_k-RIP for any rank k matrix M, $\mathcal{A}'\mathcal{A}$ is nearly isometric [Jain et al., 2012], which implies $\|M - \mathcal{A}'\mathcal{A}(M)/n\|_2 \leq \delta$. Then we can construct our estimation sequence as following:

$$\widetilde{M}^{(t)} = \frac{1}{n}\mathcal{A}'\mathcal{A}(M^* - \widetilde{M}^{(t-1)}) + \widetilde{M}^{(t-1)} \ , \ \widetilde{\boldsymbol{w}}^{(t)} = (I - \frac{1}{n}XX^\top)(\boldsymbol{w}^* - \widetilde{\boldsymbol{w}}^{(t-1)}) + \widetilde{\boldsymbol{w}}^{(t-1)} \ .$$

However, in gFM and symmetric rank-one matrix sensing, the ℓ_2-norm RIP condition cannot be satisfied with high probability [Cai and Zhang, 2015]. To establish an RIP-like condition for rank-one matrix sensing, several variants have been proposed, such as the ℓ_2/ℓ_1-RIP condition [Cai and Zhang, 2015, Chen et al., 2015]. The essential idea of these variants is to replace the ℓ_2-norm $\|\mathcal{A}(M)\|_2$ with ℓ_1-norm $\|\mathcal{A}(M)\|_1$ then a similar norm inequality can be established for all low rank matrix again. However, even using these ℓ_1-norm RIP variants, we are still unable to design an efficient alternating algorithm. All these ℓ_1-norm RIP variants have to deal with trace norm programming problems. In fact, it is impossible to construct an estimation sequence based on ℓ_1-norm RIP because we require ℓ_2-norm bound on $\mathcal{A}'\mathcal{A}$ during the construction.

A key ingredient of our framework is to propose a novel ℓ_2-norm RIP condition to overcome the above difficulty. The main technique reason for the failure of conventional ℓ_2-norm RIP is that it tries to bound $\mathcal{A}'\mathcal{A}(M)$ over all rank k matrices. This is too aggressive to be successful in rank-one matrix sensing. Regarding to our estimation sequence, what we really need is to make the RIP hold for current low rank matrix $M^{(t)}$. Once we update our estimation $M^{(t+1)}$, we can regenerate a new sensing operator independent of $M^{(t)}$ to avoid bounding $\mathcal{A}'\mathcal{A}$ over all rank k matrices. To this end, we propose the Conditionally Independent RIP (CI-RIP) condition.

Definition 3 (CI-RIP)**.** A matrix sensing operator $\mathcal{A}$ is Conditionally Independent RIP with constant δ_k, if for a fixed rank k matrix M, $\mathcal{A}$ is sampled independently regarding to M and satisfies

$$\|(I - \frac{1}{n}\mathcal{A}'\mathcal{A})M\|_2^2 \leq \delta_k \ . \tag{4}$$

An ℓ_2-norm or ℓ_1-norm RIP sensing operator is naturally CI-RIP but the reverse is not true. In CI-RIP, $\mathcal{A}$ is no longer a fixed but random sensing operator independent of M. In one-pass algorithm, this is achievable if we always retrieve new instances to construct $\mathcal{A}$ in one mini-batch updating. Usually Eq. (4) doesn't hold in a batch method since $M^{(t+1)}$ depends on $\mathcal{A}(M^{(t)})$.

An asymmetric rank-one matrix sensing operator is clearly CI-RIP due to the independency between left/right design vectors. But a symmetric rank-one matrix sensing operator is not CI-RIP. In fact it is a biased estimator since $E(\boldsymbol{x}^\top M\boldsymbol{x}) = \mathrm{tr}(M)$. To this end, we propose a shifted version of CI-RIP for symmetric rank-one matrix sensing operator in the following theorem. This theorem is the key tool in our analysis.

Theorem 4 (Shifted CI-RIP)**.** *Suppose $\boldsymbol{x}_i$ are independent standard random Gaussian vectors, M is a fixed symmetric rank k matrix independent of $\boldsymbol{x}_i$ and $\boldsymbol{w}$ is a fixed vector. Then with a probability at least $1-\eta$, provided $n \geq Ck^3d/\delta^2$,*

$$\|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M) - \frac{1}{2}\mathrm{tr}(M)I - M\|_2 \leq \delta\|M\|_2 \ .$$

Theorem 4 shows that $\frac{1}{2n}\mathcal{A}'\mathcal{A}(M)$ is nearly isometric after shifting by its expectation $\frac{1}{2}\mathrm{tr}(M)I$. The RIP constant $\delta = O(\sqrt{k^3d/n})$. In gFM, we choose $M = M^* - M^{(t)}$ therefore M is of rank $3k$.

Under the same settings of Theorem 4, suppose that $d \geq C$ then the following lemmas hold true with a probability at least $1-\eta$ for fixed $\boldsymbol{w}$ and M .

Lemma 5. $|\frac{1}{n}\mathbf{1}^\top\mathcal{A}(M)) - \mathrm{tr}(M)| \leq \delta\|M\|_2$ *provided* $n \geq Ck/\delta^2$.

Lemma 6. $|\frac{1}{n}\mathbf{1}^\top X^\top\boldsymbol{w}| \leq \|\boldsymbol{w}\|_2\delta$ *provided* $n \geq C/\delta^2$.

Lemma 7. $\|\frac{1}{n}\mathcal{A}'(X^\top\boldsymbol{w})\|_2 \leq \|\boldsymbol{w}\|_2\delta$ *provided* $n \geq Cd/\delta^2$.

Lemma 8. $\|\frac{1}{n}X^\top\mathcal{A}(M)\|_2 \leq \|M\|_2\delta$ *provided* $n \geq Ck^2d/\delta^2$.

Lemma 9. $\|I - \frac{1}{n}XX^\top\|_2 \leq \delta$ *provided* $n \geq Cd/\delta^2$.

Equipping with the above lemmas, we construct our estimation sequence as following.

Lemma 10. *Let $M^{(t)}, H_1^{(t)}, h_2^{(t)}, \boldsymbol{h}_3^{(t)}$ be defined as in Algorithm 1. Define $\epsilon_t = \|\boldsymbol{w}^* - \boldsymbol{w}^{(t)}\|_2 + \|M^* - M^{(t)}\|_2$. Then with a probability at least $1-\eta$, provided $n \geq Ck^3d/\delta^2$,*

$$H_1^{(t)} = M^* - M^{(t)} + \mathrm{tr}(M^* - M^{(t)})I + O(\delta\epsilon_t) \ , \ h_2^{(t)} = \mathrm{tr}(M^* - M^{(t)}) + O(\delta\epsilon_t)$$
$$\boldsymbol{h}_3^{(t)} = \boldsymbol{w}^* - \boldsymbol{w}^{(t)} + O(\delta\epsilon_t) \ .$$

Suppose by construction, $\epsilon_t \to 0$ when $t \to \infty$. Then $H_1^{(t)} - h_2^{(t)} I + M^{(t)} \to M^*$ and $\boldsymbol{h}_3^{(t)} + \boldsymbol{w}^{(t)} \to \boldsymbol{w}^*$ and then the proof of Theorem 1 is completed. In the following we only need to show that Lemma 10 constructs an estimation sequence with $\epsilon_t = O(\delta^t) \to 0$. To this end, we need a few things from matrix perturbation theory.

By Theorem 1, $U^{(t)}$ will converge to U^* up to column order perturbation. We use the largest canonical angle to measure the subspace distance spanned by $U^{(t)}$ and U^*, which is denoted as $\theta_t = \theta(U^{(t)}, U^*)$. For any matrix U, it is well known [Zhu and Knyazev, 2013] that

$$\sin\theta(U, U^*) = \|U_\perp^{*\top} U\|_2,\ \cos\theta(U, U^*) = \sigma_k\{U^{*\top} U\},\ \tan\theta(U, U^*) = \|U_\perp^{*\top} U (U^{*\top} U)^{-1}\|_2\ .$$

The last tangent equality allows us to bound the canonical angle after QR decomposition. Suppose $U^{(t)} R = \hat{U}^{(t)}$ in the QR step of Algorithm 1, we have

$$\begin{aligned}\tan\theta(\hat{U}^{(t)}, U^*) &= \|U_\perp^{*\top} \hat{U}^{(t)} (U^{*\top} \hat{U}^{(t)})^{-1}\|_2 = \|U_\perp^{*\top} U^{(t)} R (U^{*\top} U^{(t)} R)^{-1}\|_2 \\ &= \|U_\perp^{*\top} U^{(t)} (U^{*\top} U^{(t)})^{-1}\|_2 = \tan\theta(U^{(t)}, U^*)\ .\end{aligned}$$

Therefore, it is more convenient to measure the subspace distance by tangent function.

To show $\epsilon_t \to 0$, we recursively define the following variables:

$$\alpha_t \triangleq \tan\theta_t,\ \beta_t \triangleq \|\boldsymbol{w}^* - \boldsymbol{w}^{(t)}\|_2,\ \gamma_t \triangleq \|M^* - M^{(t)}\|_2,\ \epsilon_t \triangleq \beta_t + \gamma_t\ .$$

The following lemma derives the recursive inequalities regarding to $\{\alpha_t, \beta_t, \gamma_t\}$.

Lemma 11. *Under the same settings of Theorem 1, suppose $\alpha_t \le 2$, $\delta\epsilon_t \le 4\sqrt{5}\sigma_k^*$, then*

$$\alpha_{t+1} \le 4\sqrt{5}\delta\sigma_k^{*-1}(\beta_t + \gamma_t),\ \beta_{t+1} \le \delta(\beta_t + \gamma_t),\ \gamma_{t+1} \le \alpha_{t+1}\|M^*\|_2 + 2\delta(\beta_t + \gamma_t)\ .$$

In Lemma 11, when we choose n such that $\delta = O(1/\sqrt{n})$ is small enough, $\{\alpha_t, \beta_t, \gamma_t\}$ will converge to zero. The only question is the initial value $\{\alpha_0, \beta_0, \gamma_0\}$. According to the initialization step of gFM, $\beta_0 \le \|\boldsymbol{w}^*\|_2$ and $\gamma_0 \le \|M^*\|_2$. To bound α_0 , we need the following lemma which directly follows Wely's and Wedin's theorems [Stewart and Sun, 1990].

Lemma 12. *Denote U and $\widetilde{U}$ as the top-k left singular vectors of M and $\widetilde{M} = M + O(\epsilon)$ respectively. The i-th singular value of M is σ_i. Suppose that $\epsilon \le \frac{\sigma_k - \sigma_{k+1}}{4}$. Then the largest canonical angle between U and $\widetilde{U}$, denoted as $\theta(U, \widetilde{U})$, is bounded by $\sin\theta(U, \widetilde{U}) \le 2\epsilon/(\sigma_k - \sigma_{k+1})$.*

According to Lemma 12, when $2\delta(\|\boldsymbol{w}^*\|_2 + \|M^*\|_2) \le \sigma_k^*/4$, we have $\sin\theta_0 \le 4\delta(\|\boldsymbol{w}^*\|_2 + \|M^*\|_2)/\sigma_k^*$. Therefore, $\alpha_0 \le 2$ provided $\delta \le \sigma_k^*/[8(\|\boldsymbol{w}^*\|_2 + \|M^*\|_2)]$.

Proof of Theorem 1. Suppose that at step t, $\alpha_t \leq 2$, $\delta\epsilon_t \leq 4\sqrt{5}\sigma_k^*$, from Lemma 11,

$$\begin{aligned}\beta_{t+1} + \gamma_{t+1} \leq & \beta_{t+1} + \alpha_{t+1}\|M^*\|_2 + 2\delta(\beta_t + \gamma_t) \leq \delta\epsilon_t + 4\sqrt{5}\delta\sigma_k^{*-1}\epsilon_t\|M^*\|_2 + 2\delta\epsilon_t \\ = & (4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)\delta\epsilon_t \ .\end{aligned}$$

Therefore,

$$\begin{aligned}\epsilon_t = \beta_t + \gamma_t &\leq [(4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)\delta]^t(\beta_0 + \gamma_0) \\ \alpha_{t+1} &\leq 4\sqrt{5}\delta\sigma_k^{*-1}(\beta_t + \gamma_t) \leq 4\sqrt{5}\delta\sigma_k^{*-1}[(4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)\delta]^t(\beta_0 + \gamma_0) \ .\end{aligned}$$

Clearly we need $(4\sqrt{5}\sigma_1^*/\sigma_k^*+3)\delta < 1$ to ensure convergence, which is guaranteed by $\delta < \frac{\sigma_k^*}{4\sqrt{5}\sigma_1^*+3\sigma_k^*}$. To ensure the recursive inequality holds for any t, we require $\alpha_{t+1} \leq 2$, which is guaranteed by

$$4\sqrt{5}(\beta_0 + \gamma_0)\delta/\sigma_k^* \leq 2 \Leftrightarrow \delta \leq \frac{\sigma_k^*}{2\sqrt{5}(\sigma_1^* + \beta_0)} \ .$$

To ensure the condition $\delta\epsilon_t \leq 4\sqrt{5}\sigma_k^*$,

$$\delta \leq 4\sqrt{5}\sigma_k^*/\epsilon_0 = 4\sqrt{5}\sigma_k^*/(\sigma_1^* + \beta_0) \Rightarrow \delta \leq 4\sqrt{5}\sigma_k^*/\epsilon_t \ .$$

In summary, when

$$\begin{aligned}&\delta \leq \min\left\{\frac{\sigma_k^*}{4\sqrt{5}(\sigma_1^* + \beta_0)}, \frac{\sigma_k^*}{4\sqrt{5}\sigma_1^* + 3\sigma_k^*}, \frac{\sigma_k^*}{2\sqrt{5}(\sigma_1^* + \beta_0)}, \frac{\sigma_k^*}{8(\sigma_1^* + \beta_0)}\right\} \\ \Leftarrow &\delta \leq \frac{\sigma_k^*}{4\sqrt{5}\sigma_1^* + 3\sigma_k^* + 4\sqrt{5}\beta_0} \ .\end{aligned}$$

we have

$$\epsilon_t = [(4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)\delta]^t(\sigma_1^* + \gamma_0) \ .$$

To simplify the result, replace δ with $\delta_1 = (4\sqrt{5}\sigma_1^*/\sigma_k^* + 3)\delta$. The proof is completed. □

4.1 Noisy Case

In this subsection, we analyze the performance of gFM under noisy setting. Suppose that M^* is no longer low rank, $M^* = U^*\Lambda^*U^{*\top} + U_\perp^*\Lambda_\perp^*U_\perp^{*\top}$ where $\Lambda_\perp^* = \mathrm{diag}(\lambda_{k+1}, \cdots, \lambda_d)$ is the residual spectrum. Denote $M_k^* = U^*\Lambda^*U^{*\top}$ to be the best rank k approximation of M^* and $M_\perp^* = M^* - M_k^*$. The additive noise ξ_i's are independently sampled from subgaussian with proxy variance ξ.

First we generalize the above theorems and lemmas to noisy case.

Lemma 13. *Suppose that in Eq. (1) $\boldsymbol{x}_i$'s are independent standard random Gaussian vectors. M is a fixed rank k matrix. $M_\perp^* \neq \mathbf{0}$ and $\xi > 0$. Then provided $n \geq Ck^3d/\delta^2$, with a probability at least $1-\eta$,*

$$\|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M^*-M) - \frac{1}{2}\mathrm{tr}(M_k^*-M)I - (M_k^*-M)\|_2 \leq \delta\|M_k^*-M\|_2 + C\sigma_{k+1}^* d^2/\sqrt{n} \tag{5}$$

$$|\frac{1}{n}\mathbf{1}^\top\mathcal{A}(M^*-M) - \mathrm{tr}(M_k^*-M)| \leq \delta\|M_k^*-M\|_2 + C\sigma_{k+1}^* d^2/\sqrt{n} \tag{6}$$

$$\|\frac{1}{n}X^\top\mathcal{A}(M^*-M)\|_2 \leq \delta\|M_k^*-M\|_2 + C\sigma_{k+1}^* d^2/\sqrt{n} \tag{7}$$

$$\|\frac{1}{n}\mathcal{A}'(X^\top\boldsymbol{w})\|_2 \leq \delta\|\boldsymbol{w}\|_2,\ \|\frac{1}{n}\mathbf{1}^\top X^\top\boldsymbol{w}\|_2 \leq \delta\|\boldsymbol{w}\|_2\ . \tag{8}$$

Define $\gamma_t = \|M_k^* - M^{(t)}\|_2$ similar to the noise-free case. According to Lemma 13, when $\xi = 0$, for $n \geq Ck^3d/\delta^2$,

$$\begin{aligned}
H_1^{(t)} =& M_k^* - M^{(t)} + \frac{1}{2}\mathrm{tr}(M_k^* - M^{(t)})I + O(\delta\epsilon_t + C\sigma_{k+1}^* d^2/\sqrt{n}) \\
h_2^{(t)} =& \mathrm{tr}(M^* - M^{(t)}) + O(\delta\epsilon_t + C\sigma_{k+1}^* d^2/\sqrt{n}) \\
\boldsymbol{h}_3^{(t)} =& \boldsymbol{w}^* - \boldsymbol{w}^{(t)} + O(\delta\epsilon_t + C\sigma_{k+1}^* d^2/\sqrt{n})\ .
\end{aligned}$$

Define $r = C\sigma_{k+1}^* d^2/\sqrt{n}$. If $\xi > 0$, it is easy to check that the perturbation becomes $\hat{r} = r + O(\xi/\sqrt{n})$. Therefore we uniformly use r to present the perturbation term. The recursive inequalities regarding to the recovery error is constructed in Lemma 14.

Lemma 14. *Under the same settings of Lemma 13, define $\rho \triangleq 2\sigma_{k+1}^*/(\sigma_k^* + \sigma_{k+1}^*)$. Suppose that at any step i, $0 \leq i \leq t$, $\alpha_i \leq 2$. When provided $4\sqrt{5}(\delta\epsilon_t + r) \leq \sigma_k^* - \sigma_{k+1}^*$,*

$$\alpha_{t+1} \leq \rho\alpha_t + \frac{4\sqrt{5}}{\sigma_k^* + \sigma_{k+1}^*}\delta\epsilon_t + \frac{4\sqrt{5}}{\sigma_k^* + \sigma_{k+1}^*}r\ ,\ \beta_{t+1} \leq \delta\epsilon_t + r\ ,\ \gamma_{t+1} \leq \alpha_{t+1}\|M^*\|_2 + 2\delta\epsilon_t + 2r\ .$$

The solution to the recursive inequalities in Lemma 14 is non-trivial. Comparing to the inequalities in Lemma 11, α_{t+1} is bounded by α_t in noisy case. Therefore, if we simply follow Lemma 11 to construct recursive inequality about ϵ_t , we will quickly be overloaded by recursive expansion terms. The key construction of our solution is to bound the term $\alpha_t + 8\sqrt{5}/(\sigma_k^* + \sigma_{k+1}^*)\delta\epsilon_t$. The solution is given in the following theorem.

Theorem 15. *Define constants*

$$c = 4\sqrt{5}/(\sigma_k^* + \sigma_{k+1}^*)\ ,\ b = 3 + 4\sqrt{5}\sigma_1^*/(\sigma_k^* + \sigma_{k+1}^*)\ ,\ q = (1+\rho)/2\ .$$

Then for any $t \geq 0$,

$$\alpha_t + 2c\delta\epsilon_t \leq q^t\left(2 - \frac{(1+\rho)cr}{1-q}\right) + \frac{(1+\rho)cr}{1-q}\ . \tag{9}$$

provided

$$\delta \leq \min\{\frac{1-\rho}{4\rho\sigma_1^* c}, \frac{\rho}{2b}\} \ , \ (2 + c(\sigma_k^* - \sigma_{k+1}^*))\delta\epsilon_0 + r \leq (\sigma_k^* - \sigma_{k+1}^*) \tag{10}$$

$$4\sqrt{5}\left(4 + 2c(\sigma_k^* - \sigma_{k+1}^*)\right)\delta\epsilon_0 + 4\sqrt{5}\left(4 + (\sigma_k^* - \sigma_{k+1}^*)\right)r \leq (\sigma_k^* - \sigma_{k+1}^*)^2 \ .$$

Theorem 15 gives the convergence rate of gFM under noisy settings. We bound $\alpha_t + 2c\delta\epsilon_t$ as the index of recovery error, whose convergence rate is linear. The convergence rate is controlled by q, a constant depends on the eigen gap $\sigma_{k+1}^*/\sigma_k^*$. The final recovery error is bounded by $O(r/(1-q))$. Eq. (10) is the small noise condition to ensure the noisy recovery is possible. Generally speaking, learning a $d \times d$ matrix with $O(d)$ samples is an ill-conditioned problem when the target matrix is full rank. The small noise condition given by Eq. (10) essentially says that M^* can be slightly deviated from low rank manifold and the noise shouldn't be too large to blur the spectrum of M^*. When the noise is large, Eq. (10) will be satisfied with $n = O(d^2)$ which is the information-theoretical lower bound for recovering a full rank matrix.

5 Conclusion

In this paper, we propose a provable efficient algorithm to solve generalized Factorization Machine (gFM) and rank-one matrix sensing. Our method is based on an one-pass alternating updating framework. The proposed algorithm is able to learn gFM within $O(kd)$ memory on steaming data, has linear convergence rate and only requires matrix-vector product implementation. The algorithm takes no more than $O(k^3 d \log(1/\epsilon))$ instances to achieve $O(\epsilon)$ recovery error.

References

Mathieu Blondel, Masakazu Ishihata, Akinori Fujino, and Naonori Ueda. Polynomial Networks and Factorization Machines: New Insights and Efficient Training Algorithms. pages 850–858, 2016.

T. Tony Cai and Anru Zhang. ROP: Matrix recovery via rank-one projections. *The Annals of Statistics*, 43(1):102–138, 2015.

Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

Emmanuel J. Candes, Yonina Eldar, Thomas Strohmer, and Vlad Voroninski. Phase Retrieval via Matrix Completion. *arXiv:1109.0573*, 2011.

Yuxin Chen, Yuejie Chi, and Andrea J. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *Information Theory, IEEE Transactions on*, 61(7):4034–4059, 2015.

Mark A. Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *arXiv:1601.06422*, 2016.

Moritz Hardt. Understanding Alternating Minimization for Matrix Completion. *arXiv:1312.0925*, 2013.

Moritz Hardt and Eric Price. The Noisy Power Method: A Meta Algorithm with Applications. *arXiv:1311.2495*, 2013.

Moritz Hardt and Mary Wootters. Fast matrix completion without the condition number. *arXiv:1407.4070*, 2014.

Liangjie Hong, Aziz S. Doumith, and Brian D. Davison. Co-factorization Machines: Modeling User Interests and Predicting Individual Decisions in Twitter. In *WSDM*, pages 557–566, 2013.

Prateek Jain and Inderjit S. Dhillon. Provable inductive matrix completion. *arXiv:1306.0626*, 2013.

Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank Matrix Completion using Alternating Minimization. *arXiv:1212.0467*, 2012.

Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank Matrix Completion Using Alternating Minimization. In *STOC*, pages 665–674, 2013.

V. Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*, volume 2033. Springer, 2011.

Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Low rank matrix recovery from rank one measurements. *arXiv:1410.6913*, 2014.

Kiryung Lee, Yihong Wu, and Yoram Bresler. Near Optimal Compressed Sensing of Sparse Rank-One Matrices via Sparse Power Factorization. *arXiv:1312.0525*, 2013.

Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi. Phase Retrieval using Alternating Minimization. *arXiv:1306.0160*, 2013.

Steffen Rendle. Factorization machines. In *ICDM*, pages 995–1000, 2010.

Steffen Rendle, Zeno Gantner, Christoph Freudenthaler, and Lars Schmidt-Thieme. Fast Context-aware Recommendations with Factorization Machines. In *SIGIR*, pages 635–644, 2011.

G. W. Stewart and Ji-guang Sun. *Matrix Perturbation Theory*. Academic Press, 1990.

Joel A. Tropp. Convex recovery of a structured signal from independent random linear measurements. *arXiv:1405.1102*, 2014.

Tuo Zhao, Zhaoran Wang, and Han Liu. Nonconvex Low Rank Matrix Factorization via Inexact First Order Oracle. 2015a.

Tuo Zhao, Zhaoran Wang, and Han Liu. A Nonconvex Optimization Framework for Low Rank Matrix Estimation. In *NIPS*, pages 559–567, 2015b.

Kai Zhong, Prateek Jain, and Inderjit S. Dhillon. Efficient matrix sensing using rank-1 gaussian measurements. In *Algorithmic Learning Theory*, pages 3–18, 2015.

Peizhen Zhu and Andrew V. Knyazev. Angles between subspaces and their tangents. *Journal of Numerical Mathematics*, 21(4), 2013.

A Preliminary

In this section, we present several important theorems and lemmas in our analysis.

The following concentration inequalities are well known.

Lemma 16. *Let x_i be zero-mean sub-Gaussian distribution with variance proxy σ^2. Denote $S_n = \sum_{i=1}^n a_i x_i$ for a fixed sequence $\{a_i\}$. Then*

$$\Pr(|S_n| > t) \leq 2\exp(-\frac{t^2}{2\sigma^2(\sum_{i=1}^n a_i^2)}) \ .$$

That is, with a probability at least $1-\eta$,

$$|S_n| \leq \sigma\sqrt{\sum_{i=1}^n a_i^2}\sqrt{2\log(2/\eta)} \ .$$

Corollary 17. *Let $x_i \sim \mathcal{N}(0,1)$ be a standard Gaussian distribution. Then with a probability at least $1-\eta$,*

$$\sum_{i=1}^n a_i(x_i^2-1) \leq 2\sqrt{\sum_{i=1}^n a_i}\sqrt{2\log(2/\eta)} \ .$$

For random matrix, we have matrix concentration inequalities [**?**].

Theorem 18 (Matrix Bernstein's Inequality [**?**])**.** *Suppose $\{S_i\}_{i=1}^n$ are set of independent random matrices of dimension $d_1 \times d_2$,*

$$\|S_i - ES_i\| \leq L \ .$$

Define

$$Z = \sum_{i=1}^n S_i, \ \sigma^2 = \frac{1}{n}\max(E\|(Z-EZ)(Z-EZ)^\top\|_2, E\|(Z-EZ)^\top(Z-EZ)\|_2) \ .$$

The with a probability at least $1-\delta$, for any $0 < \epsilon < 1$,

$$\frac{1}{n}\|Z-EZ\|_2 \leq 9\epsilon\sqrt{\log((d_1+d_2)/\delta)}$$

provided

$$n \geq \max(\sigma^2, L)/\epsilon^2 \ .$$

And for any n,

$$\frac{1}{n}\|Z-EZ\|_2 \leq \frac{4}{3}\frac{L}{n}\log((d_1+d_2)/\delta) + 3\sqrt{2\frac{\sigma^2}{n}\log((d_1+d_2)/\delta)} \ .$$

Using matrix Bernstein's inequality, we can bound the covariance estimator.

Corollary 19 (Matrix Bernstein's Inequality for Covariance Estimator [**?**])**.** *Suppose $\boldsymbol{x}_i \in \mathbb{R}^d, i = 1, 2, \cdots, n$ are independent random variables with zero mean.*

$$\|\boldsymbol{x}_i\|^2 \leq B, \ A = E(\boldsymbol{x}_i\boldsymbol{x}_i{}^\top)$$

Then with a probability at least $1-\delta$,

$$\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i^\top - A\|_2 \leq 9\epsilon\sqrt{\log(2d/\delta)/n}$$

provided

$$n \geq \max(B\|A\|, B)/\epsilon^2 \ .$$

B Proof of Lemmas

B.1 Proof of Lemma 5

Proof. Denote the eigenvalue decomposition of M as

$$M = U\Lambda U^\top = U\mathrm{diag}(\lambda_1, \lambda_2, \cdots, \lambda_k)U^\top$$

Since Gaussian distribution is rotation invariant, $\hat{\boldsymbol{x}}_i = U^\top \boldsymbol{x}_i$ also follows standard Gaussian distribution.

$$\boldsymbol{x}_i{}^\top M \boldsymbol{x}_i = \boldsymbol{x}_i{}^\top U\Lambda U^\top \boldsymbol{x}_i = |\hat{\boldsymbol{x}}_i{}^\top \Lambda \hat{\boldsymbol{x}}_i| = \sum_{j=1}^{k} \lambda_j \hat{\boldsymbol{x}}_{i,j}^2 \ .$$

It is easy to see that $E(\boldsymbol{x}_i{}^\top M \boldsymbol{x}_i) = \sum_j \lambda_j = \mathrm{tr}(M)$. Define

$$a_i \triangleq \boldsymbol{x}_i{}^\top M \boldsymbol{x}_i - \mathrm{tr}(M) = \sum_{j=1}^{d} \lambda_j (\hat{\boldsymbol{x}}_{i,j}^2 - 1)$$

According to Corollary 17, for a fixed i, with a probability at least $1-\eta$,

$$|a_i| \le 2\|M\|_F \sqrt{2\log(2/\eta)} \ .$$

Then for any i, with a probability at least $1-\eta$,

$$|a_i| \le 2\|M\|_F \sqrt{2\log(2n/\eta)} \ .$$

Apply Corollary 17 again, with a probability at least $1-2\eta$,

$$\begin{aligned} |\frac{1}{n}\sum_{i=1}^{n} a_i - \mathrm{tr}(M)| \le & 2\|M\|_F \sqrt{2\log(2n/\eta)} \sqrt{2\log(2/\eta)/n} \\ \le & 2\sqrt{k}\|M\|_2 \sqrt{2\log(2n/\eta)} \sqrt{2\log(2/\eta)/n} \ . \end{aligned}$$

Denote $\delta = 2\sqrt{k}\sqrt{2\log(2n/\eta)}\sqrt{2\log(2/\eta)/n}$. Then when $n \ge Ck/\delta^2$,

$$|\frac{1}{n}\sum_{i=1}^{n} a_i - \mathrm{tr}(M)| \le \|M\|_2 \delta \ .$$

□

B.2 Proof of Lemma 6

Proof. Define random variable

$$\begin{aligned} a_i = & \boldsymbol{x}_i{}^\top \boldsymbol{w}, \ Ea_i = 0 \\ Ea_i^2 \le & \|\boldsymbol{w}\|_2^2 \end{aligned}$$

Then according to Lemma 16, with a probability at least $1-\eta$,

$$|\frac{1}{n}\sum_{i=1}^{n} a_i| \le \|\boldsymbol{w}\|_2 \sqrt{2\log(2/\eta)/n} \ .$$

□

B.3 Proof of Lemma 8

Proof. Define random vector

$$\boldsymbol{a}_i = \boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i,\ E\boldsymbol{a}_i = 0\ .$$

With a probability at least $(1-\eta_1)(1-\eta_2)$,

$$\begin{aligned}\|\boldsymbol{a}_i\|_2 =& \|\boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i\|_2 \leq \|\boldsymbol{x}_i^\top M \boldsymbol{x}_i\|_2 \|\boldsymbol{x}_i\|_2 \\ \leq& (|\mathrm{tr}(M)| + 2\|M\|_F \sqrt{2\log(2n/\eta_1)}) \sqrt{2d\log(2n/\eta_2)} \\ \triangleq& c_1 \sqrt{2d\log(2n/\eta_2)}\ .\end{aligned}$$

$$\begin{aligned}\|E\boldsymbol{a}_i^\top \boldsymbol{a}_i\|_2 =& \|\boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top \boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i\|_2 \\ \leq& (\boldsymbol{x}_i^\top M \boldsymbol{x}_i)^2 \|\boldsymbol{x}_i\|_2^2 \\ \leq& 2c_1^2 d\log(2n/\eta_2)\ .\end{aligned}$$

By matrix Bernstein's inequality, the concentration holds when

$$\begin{aligned}n \geq& \frac{1}{\epsilon^2} \max\{c_1\sqrt{2d\log(2n/\eta_2)}, 2c_1^2 d\log(2n/\eta_2)\} \\ =& \frac{1}{\epsilon^2} O(k^2 d\|M\|_2^2)\ .\end{aligned}$$

Therefore, after taking the union bound, there exists some constant $C_2 = O(\log(2n/\eta))$,

$$\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{a_i}\|_2 \leq \epsilon$$

$$n \geq C_2 k^2 d\|M\|_2^2 \log(2(d+1)/\eta)/\epsilon^2\ .$$

Denote $\delta = \|M\|_2/\epsilon$. Then when $n \geq Ck^2 d/\delta$,

$$\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{a_i}\|_2 \leq \|M\|_2 \delta\ .$$

□

B.4 Proof of Lemma 7

$$\|\frac{1}{n}\mathcal{A}'(X^\top \boldsymbol{w})\|_2 = \|\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i \boldsymbol{x}_i^\top \boldsymbol{w} \boldsymbol{x}_i^\top\|_2\ .$$

$$E\{\boldsymbol{x}_i \boldsymbol{x}_i^\top \boldsymbol{w} \boldsymbol{x}_i^\top\} = 0$$

$$\begin{aligned}\|\boldsymbol{x}_i \boldsymbol{x}_i^\top \boldsymbol{w} \boldsymbol{x}_i^\top\|_2 \leq& \|\boldsymbol{x}_i^\top \boldsymbol{w}\|_2 \|\boldsymbol{x}_i\|_2^2 \\ \leq& 2\|\boldsymbol{w}\|_2 \sqrt{2\log(2/\eta)}(d + 2\sqrt{2d\log(2n/\eta)}) \\ \leq& 4\|\boldsymbol{w}\|_2 \sqrt{2\log(2/\eta)} d\end{aligned}$$

provided $d \geq 8\log(2n/\eta)$. Now considering

$$\{E\boldsymbol{x}_i\boldsymbol{x}_i{}^\top \boldsymbol{w}\boldsymbol{x}_i{}^\top \boldsymbol{x}_i\boldsymbol{w}^\top \boldsymbol{x}_i\boldsymbol{x}_i{}^\top\}_{pq} = E\{(\sum x_p x_q w_{i1}x_{i1}w_{i2}x_{i2}x_{i3}^2)\}$$

When $p \neq q$,

$$\begin{aligned}
E\{(\sum x_p x_q w_{i1}x_{i1}w_{i2}x_{i2}x_{i3}^2)\} &= E\{(2\sum_{i3} x_p x_q w_p x_p w_q x_q x_{i3}^2)\} \\
&= E\{(2\sum_{i3} x_p^2 x_q^2 w_p w_q x_{i3}^2)\} \\
&= 2w_p w_q E\{(\sum_{i3} x_p^2 x_q^2 x_{i3}^2)\} \\
&= 2w_p w_q d
\end{aligned}$$

When $p = q$,

$$\begin{aligned}
E\{(\sum x_p x_q w_{i1}x_{i1}w_{i2}x_{i2}x_{i3}^2)\} =& E\{(\sum x_p^2 w_{i1}x_{i1}w_{i2}x_{i2}x_{i3}^2)\} \\
=& E\{(\sum x_p^2 w_p x_p w_p x_p x_{i3}^2 + \sum x_p^2 w_j x_j w_j x_j x_{i3}^2 + \sum x_p^2 w_{i3}x_{i3}w_{i3}x_{i3}x_{i3}^2)\} \\
=& E\{(\sum_{i3\neq p} x_p^4 w_p^2 x_{i3}^2 + \sum_{j\neq i3\neq p} x_p^2 w_j^2 x_j^2 x_{i3}^2 + \sum_{i3\neq p} x_p^2 w_{i3}^2 x_{i3}^4)\} \\
=& w_p^2(d-1) + \sum_{j\neq p} w_j^2(d-1) + \sum_{i3\neq p} w_{i3}^2 \\
=& w_p^2(d-1) + \sum_{j\neq p} w_j^2 d = w_p^2(d-1) + \sum_{j=1}^d w_j^2 d - w_p^2 d \\
=& \sum_{j=1}^d w_j^2 d - w_p^2
\end{aligned}$$

Therefore,

$$E\boldsymbol{x}_i\boldsymbol{x}_i{}^\top \boldsymbol{w}\boldsymbol{x}_i{}^\top \boldsymbol{x}_i\boldsymbol{w}^\top \boldsymbol{x}_i\boldsymbol{x}_i{}^\top = d\mathrm{diag}\{\|\boldsymbol{w}\|_2^2\} - \mathrm{diag}\{\boldsymbol{w}\circ\boldsymbol{w}\} + 2d\boldsymbol{w}\boldsymbol{w}^\top$$

$$\|E\boldsymbol{x}_i\boldsymbol{x}_i{}^\top \boldsymbol{w}\boldsymbol{x}_i{}^\top \boldsymbol{x}_i\boldsymbol{w}^\top \boldsymbol{x}_i\boldsymbol{x}_i{}^\top\|_2 \leq 4d\|\boldsymbol{w}\|_2^2$$

Using matrix Bernstein's inequality,

$$\begin{aligned}
\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i{}^\top \boldsymbol{w}\boldsymbol{x}_i{}^\top\|_2 \leq& \frac{4}{3}\frac{4\|\boldsymbol{w}\|_2\sqrt{2\log(2/\eta)d}}{n}\log(2d/\eta) \\
&+ 3\sqrt{2\frac{4d\|\boldsymbol{w}\|_2^2}{n}\log(2d/\eta)} \\
\leq& C\|\boldsymbol{w}\|_2\sqrt{\frac{d}{n}}
\end{aligned}$$

Denote $\delta = C\sqrt{d/n}$, when $n \geq Cd/\delta^2$, $d \geq 8\log(2n/\eta)$,

$$\|\frac{1}{n}\sum_{i=1}^n \boldsymbol{x}_i\boldsymbol{x}_i{}^\top \boldsymbol{w}\boldsymbol{x}_i{}^\top\|_2 \leq \|\boldsymbol{w}\|_2\delta$$

B.5 Proof of Lemma 9

According to Corollary 19, when $d \geq 8\log(2n/\eta)$,

$$\|\boldsymbol{x}_i\|^2 \leq 2d$$

Therefore, with a probability at least $1-\eta$,

$$\|I - \frac{1}{n}XX^\top\|_2 \leq 9\epsilon\sqrt{\log(2d/\eta)/n}$$

for $n \geq 2d/\epsilon^2$. Denote $\delta = 9\epsilon\sqrt{\log(2d/\eta)/n}$, then when $n \geq Cd/\delta^2$,

$$\|I - \frac{1}{n}XX^\top\|_2 \leq \delta \ .$$

B.6 Proof of Lemma 11

To derive α_{t+1} ,

$$\begin{aligned}\|U_\perp^{*\top}[M^* + O(2\delta\epsilon_t)]U^{(t)}\|_2 \leq& \|U_\perp^{*\top}M^*U^{(t)}\|_2 + 2\delta\epsilon_t \\ \leq& 2\delta\epsilon_t\end{aligned}$$

$$\begin{aligned}\sigma_k\{U^{*\top}[M^* + O(2\delta\epsilon_t)]U^{(t)}\} \geq& U^{*\top}M^*U^{(t)} - 2\delta\epsilon_t \\ \geq& \sigma_k^*\sigma_k\{U^{*\top}U^{(t)}\} - 2\delta\epsilon_t \\ =& \sigma_k^*\cos\theta_t - 2\delta\epsilon_t\end{aligned}$$

$$\begin{aligned}\alpha_{t+1} = \tan\theta_{t+1} =& \frac{\|U_\perp^{*\top}[M^* + O(2\delta\epsilon_t)]U^{(t)}\|_2}{\sigma_k\{U^{*\top}[M^* + O(2\delta\epsilon_t)]U^{(t)}\}} \\ \leq& \frac{2\delta\epsilon_t}{\sigma_k^*\cos\theta_t - 2\delta\epsilon_t} \ .\end{aligned}$$

According to the assumption, $\cos\theta_t \geq \frac{1}{\sqrt{5}}$, $2\delta\epsilon_t \leq \frac{1}{2\sqrt{5}}\sigma_k^*$, therefore

$$\alpha_{t+1} \leq 2\sqrt{5}\epsilon_t/\sigma_k^* = 4\sqrt{5}\delta(\beta_t + \gamma_t)/\sigma_k^* \ .$$

To derive γ_{t+1},

$$\begin{aligned}\gamma_{t+1} =& \|M^* - M^{(t+1)}\|_2 \\ =& \|M^* - (U^{(t+1)}U^{(t+1)\top}(H_1^{(t)} - H_2^{(t)} + M^{(t)})^\top)\|_2 \\ =& \|M^* - U^{(t+1)}U^{(t+1)\top}(M^* + O(2\delta(\gamma_t + \beta_t)))^\top\|_2 \\ =& \|(I - U^{(t+1)}U^{(t+1)\top})M^* + U^{(t+1)}U^{(t+1)\top}O(2\delta(\gamma_t + \beta_t)))^\top\|_2 \\ \leq& \|(I - U^{(t+1)}U^{(t+1)\top})M^*\|_2 + O(2\delta(\gamma_t + \beta_t)) \\ \leq& \tan\theta_{t+1}\|M^*\|_2 + 2\delta(\gamma_t + \beta_t) \\ =& \alpha_{t+1}\|M^*\|_2 + 2\delta(\gamma_t + \beta_t) \ .\end{aligned}$$

The recursive inequality of β_t is trivial.

C Proof of Theorem 4

Proof. Denote $\sigma_1 = \|M\|_2$. Define random matrix

$$B_i = \boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top \ .$$

It is easy to check that

$$EB_i = 2M + \mathrm{tr}(M) I \ .$$

$$\begin{aligned}
\|B_i - EB_i\|_2 =& \|\boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top - 2M - \mathrm{tr}(M) I\|_2 \\
\leq& \|\boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top\|_2 + \|2M - \mathrm{tr}(M) I\|_2 \\
\leq& \|\boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top\|_2 + 2\|M\|_2 + |\mathrm{tr}(M)| \ .
\end{aligned}$$

According to Lemma 5, with a probability at least $1 - \eta_2$, for any $i \in \{1, \cdots, n\}$,

$$|\boldsymbol{x}_i^\top M \boldsymbol{x}_i| \leq |\mathrm{tr}(M)| + 2\|M\|_F \sqrt{2\log(2n/\eta_2)} \triangleq c_1 \ .$$

Therefore we have, with a probability at least $(1 - \eta_1)(1 - \eta_2)$,

$$\begin{aligned}
\|B_i - EB_i\|_2 \leq& \|\boldsymbol{x}_i\|_2^2 \ |\boldsymbol{x}_i^\top M \boldsymbol{x}_i| + 2\|M\|_2 + |\mathrm{tr}(M)| \\
\leq& 2d\log(2n/\eta_1)|\mathrm{tr}(M)| + 2\|M\|_F \sqrt{2\log(2n/\eta_2)} + 2\|M\|_2 + |\mathrm{tr}(M)| \\
\leq& Cdk\sigma_1 \ .
\end{aligned}$$

Next we need to bound

$$\begin{aligned}
\|E(B_i - EB_i)(B_i - EB_i)^\top\|_2 = \|E(B_i^2) - (EB_i)^2\|_2 \leq& \|E(B_i^2)\|_2 + \|EB_i\|_2^2 \\
\leq& \|E(B_i^2)\|_2 + 2|\mathrm{tr}(M)|^2 + 2\|M\|_2^2
\end{aligned}$$

To bound $\|E(B_i^2)\|_2$, denote the eigenvalue decomposition of M as

$$M = U\Lambda U^\top = U \mathrm{diag}(\lambda_1, \lambda_2, \cdots, \lambda_k) U^\top$$

Let $U_\perp$ be the complementary basis matrix of U. Define random variables $\boldsymbol{u}_i \triangleq U^\top \boldsymbol{x}_i$, $\boldsymbol{v}_i \triangleq U_\perp^\top \boldsymbol{x}_i$. Since $\boldsymbol{x}_i$ are standard random Gaussian, $\boldsymbol{u}$ and $\boldsymbol{v}$ should also be jointly random Gaussian thus independent.

$$\begin{aligned}
\|E(B_i^2)\|_2 =& \|E(\boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top \boldsymbol{x}_i \boldsymbol{x}_i^\top M \boldsymbol{x}_i \boldsymbol{x}_i^\top)\|_2 \\
=& \|E(\begin{bmatrix} \boldsymbol{u}_i \\ \boldsymbol{v}_i \end{bmatrix} \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i (\|\boldsymbol{u}_i\|_2^2 + \|\boldsymbol{v}_i\|_2^2) \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \begin{bmatrix} \boldsymbol{u}_i \\ \boldsymbol{v}_i \end{bmatrix}^\top)\|_2 \\
\leq& \|E(\boldsymbol{u}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i (\|\boldsymbol{u}_i\|_2^2 + \|\boldsymbol{v}_i\|_2^2) \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i^\top\|_2 \\
&+ 2\|E(\boldsymbol{u}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i (\|\boldsymbol{u}_i\|_2^2 + \|\boldsymbol{v}_i\|_2^2) \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i^\top\|_2 \\
&+ \|E(\boldsymbol{v}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i (\|\boldsymbol{u}_i\|_2^2 + \|\boldsymbol{v}_i\|_2^2) \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i^\top\|_2 \\
\leq& \|E(\boldsymbol{u}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i^\top)\|_2 + \|E(\boldsymbol{u}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i^\top)\|_2 \\
&+ 2\|E(\boldsymbol{u}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i^\top)\|_2 + 2\|E(\boldsymbol{u}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i^\top)\|_2 \\
&+ \|E(\boldsymbol{v}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i^\top)\|_2 + \|E(\boldsymbol{v}_i \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i^\top)\|_2 \ .
\end{aligned}$$

Let us bound the above 6 terms respectively. Recall that with a probability at least $1-\eta_2$,

$$|\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i| = |\boldsymbol{x}_i{}^\top M \boldsymbol{x}_i| \le c_1 \ .$$

With a probability at least $1-\eta_3$, for any $i \in \{1, \cdots, n\}$, $\|\boldsymbol{u}_i\|_2 \le 2\sqrt{k\log(4n/\eta_3)}$,$\|\boldsymbol{v}_i\|_2 \le 2\sqrt{d\log(4n/\eta_3)}$. Then:

$$\begin{aligned}
&\|E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i{}^\top)\|_2 \\
=&\|E\{\left((\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2 \|\boldsymbol{u}_i\|_2^2\right) \boldsymbol{u}_i \boldsymbol{u}_i{}^\top\}\|_2 \\
\le&(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2 \|\boldsymbol{u}_i\|_2^4 \\
\le&32 c_1^2 k^2 \log^2(2n/\eta_3) \ .
\end{aligned}$$

$$\begin{aligned}
&\|E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i{}^\top)\|_2 \\
=&\|E(\|\boldsymbol{v}_i\|_2^2) E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i{}^\top)\|_2 \\
\le&4d\log(4n/\eta_3)\|E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{u}_i{}^\top)\|_2 \\
\le&4d\log(4n/\eta_3)\|\boldsymbol{u}_i\|_2^2 (\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2 \\
\le&4d\log(4n/\eta_3) c_1^2 (4k\log(4n/\eta_3)) \ .
\end{aligned}$$

$$\begin{aligned}
&2\|E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i{}^\top)\|_2 \\
=&2\|E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i) E(\boldsymbol{v}_i{}^\top)\|_2 = 0
\end{aligned}$$

$$\begin{aligned}
&2\|E(\boldsymbol{u}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i{}^\top)\|_2 \\
=&2\|E(\boldsymbol{u}_i(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2) E(\|\boldsymbol{v}_i\|_2^2 \boldsymbol{v}_i{}^\top)\|_2 \\
=&2\|E(\boldsymbol{u}_i(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2) E(\boldsymbol{v}_i{}^\top \boldsymbol{v}_i \boldsymbol{v}_i{}^\top)\|_2 = 0
\end{aligned}$$

$$\begin{aligned}
&\|E(\boldsymbol{v}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i{}^\top)\|_2 \\
=&\|E(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i) E(\boldsymbol{v}_i \boldsymbol{v}_i{}^\top)\|_2 \\
=&\|E(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{u}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i)\|_2 \\
\le&(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2 \|\boldsymbol{u}_i\|_2^2 \\
\le&4c_1^2 k\log(4n/\eta_3) \ .
\end{aligned}$$

$$\begin{aligned}
&\|E(\boldsymbol{v}_i\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{u}_i{}^\top \Lambda \boldsymbol{u}_i \boldsymbol{v}_i{}^\top)\|_2 \\
=&\|E\{(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2\} E(\boldsymbol{v}_i \|\boldsymbol{v}_i\|_2^2 \boldsymbol{v}_i{}^\top)\|_2 \\
=&\|E\{(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2\}(d+2) I\|_2 \\
\le&(d+2)(\boldsymbol{u}_i^\top \Lambda \boldsymbol{u}_i)^2 \\
\le&c_1^2(d+2)
\end{aligned}$$

Add all above together, we have

$$\begin{aligned}\|E(B_i^2)\|_2 \leq & 32c_1^2k^2\log^2(2n/\eta_3) + 4d\log(4n/\eta_3)c_1^2(4k\log(4n/\eta_3)) \\ & + 4c_1^2k\log(4n/\eta_3) + c_1^2(d+2) \\ \leq & Ck^3d\sigma_1 \ .\end{aligned}$$

Apply matrix Bernsterin's inequality, the proof is completed. □

D Proof of Lemma 15

We assume that $n \geq Ck^3d/\delta^2$.

To prove Eq. (5)

$$\begin{aligned}& \|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M^* - M) - \frac{1}{2}\mathrm{tr}(M_k^* - M)I - (M_k^* - M)\|_2 \\ \leq & \|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M_k^* - M) - \frac{1}{2}\mathrm{tr}(M_k^* - M)I - (M_k^* - M)\|_2 + \|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M_\perp^*)\|_2 \\ \leq & \|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M_\perp^*)\|_2 + \delta\|M_k^* - M\|_2 \ .\end{aligned}$$

The last inequality is because of Theorem 4. To bound the first term in the last inequality, define random matrix

$$B_i = \boldsymbol{x}_i\boldsymbol{x}_i{}^\top M_\perp^* \boldsymbol{x}_i\boldsymbol{x}_i{}^\top$$

As proved in Theorem 4, $EB_i = 2M_\perp^* + \mathrm{tr}(M_\perp^*)I$.

$$\begin{aligned}\|(B_i - EB_i)\|_2 = & \|\boldsymbol{x}_i\boldsymbol{x}_i{}^\top M_\perp^* \boldsymbol{x}_i\boldsymbol{x}_i{}^\top - 2M_\perp^* + \mathrm{tr}(M_\perp^*)I\|_2 \\ \leq & \|\boldsymbol{x}_i\boldsymbol{x}_i{}^\top M_\perp^* \boldsymbol{x}_i\boldsymbol{x}_i{}^\top\|_2 + 2\|M_\perp^*\|_2 + \|\mathrm{tr}(M_\perp^*)I\|_2 \\ = & \|\boldsymbol{x}_i\boldsymbol{x}_i{}^\top M_\perp^* \boldsymbol{x}_i\boldsymbol{x}_i{}^\top\|_2 + 2\sigma_{k+1}^* + |\mathrm{tr}(M_\perp^*)|\end{aligned}$$

While

$$\begin{aligned}\|\boldsymbol{x}_i\boldsymbol{x}_i{}^\top M_\perp^* \boldsymbol{x}_i\boldsymbol{x}_i{}^\top\|_2 \leq & \|M_\perp^*\|_2\|\boldsymbol{x}_i\|_2^4 \\ \leq & \sigma_{k+1}^*(d + 2\sqrt{2d\log(2n/\eta)})^2 \\ \leq & Cd^2\sigma_{k+1}^*\end{aligned}$$

Applying matrix Bernstein's inequality, with a probability at least $1-\eta$, we have

$$\|\frac{1}{n}\sum_{i=1}^n (B_i - EB_i)\|_2 \leq C\sigma_{k+1}^* d^2/\sqrt{n} \ .$$

Therefore

$$\|\frac{1}{2n}\mathcal{A}'\mathcal{A}(M^* - M) - \frac{1}{2}\mathrm{tr}(M_k^* - M)I - (M_k^* - M)\|_2 \leq \delta\|M_k^* - M\|_2 + C\sigma_{k+1}^{*2}d^4/\sqrt{n} \ .$$

The other inequalities can be similarly proved.

E Proof of Lemma 14

First we bound α_{t+1}. According to assumption, when

$$2(\delta\epsilon_t + r) \leq \frac{\sigma_k^* - \sigma_{k+1}^*}{2\sigma_k^*}\sigma_k^*/\sqrt{5}$$

we have

$$\begin{aligned}
\alpha_{t+1} \leq & \frac{\sigma_{k+1}^* \sin\theta_t + 2(\delta\epsilon_t + r)}{\sigma_k^* \cos\theta_t - 2(\delta\epsilon_t + r)} \\
\leq & \frac{2\sigma_k^*}{\sigma_k^* + \sigma_{k+1}^*} \frac{\sigma_{k+1}^* \sin\theta_t + 2(\delta\epsilon_t + r)}{\sigma_k^* \cos\theta_t} \\
\leq & \frac{2\sigma_{k+1}^*}{\sigma_k^* + \sigma_{k+1}^*} \tan\theta_t + \frac{2}{\sigma_k^* + \sigma_{k+1}^*} \frac{2(\delta\epsilon_t + r)}{\cos\theta_t} \\
\leq & \frac{2\sigma_{k+1}^*}{\sigma_k^* + \sigma_{k+1}^*} \tan\theta_t + \frac{4\sqrt{5}}{\sigma_k^* + \sigma_{k+1}^*}(\delta\epsilon_t + r) \\
\leq & \rho\alpha_t + \frac{4\sqrt{5}}{\sigma_k^* + \sigma_{k+1}^*}\delta\epsilon_t + \frac{4\sqrt{5}}{\sigma_k^* + \sigma_{k+1}^*} r \ .
\end{aligned}$$

To bound β_{t+1}. Clearly $\beta_{t+1} \leq \delta\epsilon_t + r$.
To bound γ_{t+1}, following the noise-free case,

$$\gamma_{t+1} \leq \alpha_{t+1}\|M^*\|_2 + 2\delta\epsilon_t + 2r \ .$$

F Proof of Lemma 15

Abbreviate

$$c = \frac{4\sqrt{5}}{\sigma_k^* + \sigma_{k+1}^*}$$

Then

$$\alpha_{t+1} \leq \rho\alpha_t + c\delta\epsilon_t + cr \ .$$

According to Lemma 14,

$$\begin{aligned}
\beta_{t+1} + \gamma_{t+1} \leq & \delta\epsilon_t + r + \alpha_{t+1}\|M^*\|_2 + 2\delta\epsilon_t + 2r \\
= & \sigma_1^*\alpha_{t+1} + 3\delta\epsilon_t + 3r \\
\leq & \sigma_1^*(\rho\alpha_t + c\delta\epsilon_t + cr) + 3\delta\epsilon_t + 3r \\
= & \rho\sigma_1^*\alpha_t + (\sigma_1^* c + 3)\delta\epsilon_t + (\sigma_1^* c + 3)r
\end{aligned}$$

Therefore, abbreviate $b \triangleq (\sigma_1^* c + 3)$,

$$\begin{cases}
\alpha_{t+1} \leq \rho\alpha_t + c\delta\epsilon_t + cr \\
\epsilon_{t+1} \leq \rho\sigma_1^*\alpha_t + b\delta\epsilon_t + br
\end{cases}$$

define

$$f_t = \alpha_t + 2c\delta\epsilon_t$$

$$
\begin{aligned}
f_{t+1} =& a_{t+1} + 2c\delta\epsilon_{t+1} \\
\leq& \rho\alpha_t + c\delta\epsilon_t + cr + 2c\delta(\rho\sigma_1^*\alpha_t + b\delta\epsilon_t + br) \\
=& \rho\alpha_t + c\delta\epsilon_t + cr + 2c\delta\rho\sigma_1^*\alpha_t + 2c\delta b\delta\epsilon_t + 2c\delta br \\
=& (\rho + 2c\delta\rho\sigma_1^*)\alpha_t + (c + 2c\delta b)\delta\epsilon_t + (1 + 2\delta b)cr
\end{aligned}
$$

When

$$
\begin{aligned}
&\delta \leq \frac{1-\rho}{4\rho\sigma_1^* c} \\
\Rightarrow &\rho + 2c\delta\rho\sigma_1^* \leq \frac{1+\rho}{2}
\end{aligned}
$$

And when

$$
\begin{aligned}
&\Rightarrow \delta \leq \frac{\rho}{2b} \\
&\Rightarrow 2\delta b \leq \rho \\
&\Rightarrow 2c\delta b \leq \rho c \\
&\Rightarrow c + 2c\delta b \leq (1+\rho)c \\
&\Rightarrow c + 2c\delta b \leq \frac{1+\rho}{2} 2c
\end{aligned}
$$

Then abbreviate $R \triangleq (c + 2c\delta b)\delta\epsilon_t + (1 + 2\delta b)cr$ we have

$$
f_{t+1} \leq \frac{1+\rho}{2} f_t + (1 + 2\delta b)cr \leq \frac{1+\rho}{2} f_t + (1+\rho)cr
$$

Abbreviate $q = (1+\rho)/2$,

$$
f_t \leq \frac{(1+\rho)cr}{1-q} + q^t (f_0 - \frac{(1+\rho)cr}{1-q})
$$

To ensure $\alpha_{t+1} \leq 2$, we require

$$
\begin{aligned}
&f_0 \leq 2 \\
\Leftarrow &\alpha_0 + 2c\delta\epsilon_0 \leq 2
\end{aligned}
$$

According to Lemma 12,

$$
\alpha_0 \leq \frac{2}{\sigma_k^* - \sigma_{k+1}^*} 2(\delta\epsilon_0 + r) = \frac{4}{\sigma_k^* - \sigma_{k+1}^*}(\delta\epsilon_0 + r)
$$

$$
\begin{aligned}
&\alpha_0 + 2c\delta\epsilon_0 \leq 2 \\
\Leftarrow &\frac{4}{\sigma_k^* - \sigma_{k+1}^*}(\delta\epsilon_0 + r) + 2c\delta\epsilon_0 \leq 2 \\
\Leftarrow &(4 + 2c(\sigma_k^* - \sigma_{k+1}^*))\delta\epsilon_0 + 4r \leq 2(\sigma_k^* - \sigma_{k+1}^*) \\
\Leftarrow &(2 + c(\sigma_k^* - \sigma_{k+1}^*))\delta\epsilon_0 + r \leq (\sigma_k^* - \sigma_{k+1}^*)
\end{aligned}
$$

In summary,

$$\alpha_t + 2c\delta\epsilon_t \leq q^t(f_0 - \frac{(1+\rho)cr}{1-q}) + \frac{(1+\rho)cr}{1-q}$$

provided

$$\delta \leq \min\{\frac{1-\rho}{4\rho\sigma_1^* c}, \frac{\rho}{2b}\}$$

and

$$(2 + c(\sigma_k^* - \sigma_{k+1}^*))\delta\epsilon_0 + r \leq (\sigma_k^* - \sigma_{k+1}^*)$$
$$4\sqrt{5}(\delta \max_t \epsilon_t + r) \leq \sigma_k^* - \sigma_{k+1}^*$$

To ensure the last inequality,

$$\delta \max_t \epsilon_t \leq f_0 \leq \alpha_0 + 2c\delta\epsilon_0 \leq \frac{4}{\sigma_k^* - \sigma_{k+1}^*}(\delta\epsilon_0 + r) + 2c\delta\epsilon_0$$
$$= (\frac{4}{\sigma_k^* - \sigma_{k+1}^*} + 2c)\delta\epsilon_0 + \frac{4}{\sigma_k^* - \sigma_{k+1}^*}r$$

Therefore we need the condition

$$4\sqrt{5}\left(\frac{4}{\sigma_k^* - \sigma_{k+1}^*} + 2c\right)\delta\epsilon_0 + 4\sqrt{5}\left(\frac{4}{\sigma_k^* - \sigma_{k+1}^*} + 1\right)r \leq \sigma_k^* - \sigma_{k+1}^*$$
$$\Leftarrow 4\sqrt{5}\left(4 + 2c(\sigma_k^* - \sigma_{k+1}^*)\right)\delta\epsilon_0 + 4\sqrt{5}\left(4 + (\sigma_k^* - \sigma_{k+1}^*)\right)r \leq (\sigma_k^* - \sigma_{k+1}^*)^2$$